

Compression and Signaling of Quantization Matrices

R. Joshi, J. Sole and M. Karczewicz

Qualcomm

AVC Method


- Zigzag scan, DPCM, signed exponential Golomb coding
 - Achieves roughly 3 bits/sample
 - 1 bit is used for sign for non-zero prediction error residuals

Proposed Method

- Better compression achievable?
 - What if the prediction error distribution is skewed towards positive entries?
 - Exploit quantization matrix properties
- Well designed quantization matrices use properties of human visual system.
 - Contrast sensitivity function (CSF)

Contrast Sensitivity Function (CSF)

- 1-D CSF is low-pass:
 - Threshold contrast increases with increasing spatial frequency



CSF - *continued*

- Quantization matrix entries are typically increasing in the row as well as column directions.

6	9	12	22
8	12	24	31
12	16	30	35
13	17	32	38

Proposed Method

- Raster scan
- Predictor
 - Max(top,left)
 - Prediction residual mostly positive
- Unsigned Golomb code (with offset)
 - Mapping of signed values to unsigned uses an offset
- Offset and Golomb parameter signaled

Unsigned Golomb code (with offset)

Source symbol (X)	Remapped symbol (Y)
0	0
1	1
2	2
3	3
-1	4
4	5
-2	6
5	7
...	...

Results

- Combination with 3 flags proposed in JCTVC-G083
 - Downsample, sym45, sym135
- Results for both symmetric and asymmetric matrices used in CE4

Results

			Symmetric quant matrix			Asymmetric quant matrix		
Down-sample	Sym45	Sym135	Bits (AVC)	Bits (proposed)	Avg. abs error	Bits (AVC)	Bits (proposed)	Avg. abs error
0	0	0	8856	7171 (81%)	0.00	9512	6450 (68%)	0.00
0	0	1	5020	3828 (76%)	0.00	5056	3503 (69%)	3.53
0	1	0	5052	4201 (83%)	5.06	4558	3552 (78%)	3.14
0	1	1	2861	2305 (81%)	5.06	2626	2016 (77%)	5.25
1	0	0	3631	2889 (80%)	0.43	3784	2539 (67%)	0.32
1	0	1	2339	1645 (70%)	0.43	2323	1501 (65%)	3.76
1	1	0	2263	1858 (82%)	4.73	1923	1472 (77%)	3.54
1	1	1	1477	1120 (76%)	4.73	1249	955 (76%)	5.52

Signaling of Partial Quantization Matrix

- Large percentage of energy concentrated in lower frequency coefficients
- An encoder may only be concerned with quantization matrix entries for lower frequencies
 - No way to signal a partial quantization matrix

Proposed Signaling

- For each quantization matrix
 - Signal `partial_quant_matrix_flag`
 - If flag is 1,
 - Signal `cornerX` and `cornerY` using fixed number of bits equal to \log_2 of the respective block sizes.
 - `cornerX` and `cornerY` are coordinates of the bottom-right corner of the rectangle for which quantization matrix entries will be signaled.
 - Signal quantization matrix entries
 - Same method used for signaling full quantization matrix is used for signaling partial quantization matrix

Conclusions

- A method for compression of quantization matrix
 - 21% bit savings over AVC method for symmetric quantization matrix
 - 28% bit savings over AVC method for asymmetric quantization matrix
 - Can be combined with other proposals.
- Partial signaling of quantization matrix